Question-answer dynamics in deductive fallacies without language

WooJin Chung (woojin@snu.ac.kr)
Department of Linguistics, Seoul National University
1 Gwanak-ro, Gwanak-gu, 08826 Seoul, Republic of Korea

Nadine Bade (nadine.bade@uni-potsdam.de)
Department of Linguistics, Collaborative Research Center 1287, University of Potsdam
Karl-Liebknecht-Strasse 24-25, 14476 Potsdam, Germany

Sam Blanc-Cuenca (sam.blanc-cuenca@ens.fr)
Institut Jean Nicod, Département d'études cognitives,
Ecole Normale Supérieure, EHESS, CNRS, PSL University
29 rue d'Ulm, 75005 Paris, France

Salvador Mascarenhas (salvador.mascarenhas@ens.fr)
Institut Jean Nicod, Département d'études cognitives,
Ecole Normale Supérieure, EHESS, CNRS, PSL University
29 rue d'Ulm, 75005 Paris, France

Abstract

We introduce purely visual paradigms that convey the logical structure of illusory inferences from disjunction: $(a \wedge b) \vee c$, $a \vdash_? b$. Although the logical information was conveyed entirely via non-linguistic means, we found that the visual paradigms induce reasoning fallacies, though less attractive than their linguistic counterparts. The visual paradigms highlight the role of alternative-based reasoning, or question-answer dynamics, as they control for narrowly interpretive processes that confound the study of their linguistic counterparts. To our knowledge, this is the first work to develop visual paradigms that represent reasoning fallacies committed by adults and involve multiple logical operators non-trivially embedded. Previous studies focused on pre-verbal children or non-human animals, and for this reason limited the scope of research to visually representing logically simple, valid inferences.

Keywords: reasoning; illusory inferences; disjunction

Introduction

Experimental investigations into the reasoning faculties of human adults are almost invariably conducted by means of linguistically presented stimuli. For example, in their seminal article on the *conjunction fallacy*, Tversky and Kahneman (1983) presented participants with a descriptive paragraph about a person named Linda, mentioning among other things her involvement in various activist movements while in college. When asked to judge which of (A) "Linda is a bank teller" or (B) "Linda is a bank teller and she is active in the feminist movement" was more likely, participants overwhelmingly chose option (B), in apparent violation of the classical probability calculus.

Yet, when reasoning based on linguistically presented information, we expect human adults to engage in interpretive semantic and pragmatic processes, which can confound our interpretation of their reasoning and decision-making behavior (Stenning & van Lambalgen, 2008). Continuing with the same example, the conjunction "fallacy" might be no fallacy at all if participants interpret option (A) as (A′) "Linda is a bank teller and she is *not* active in the feminist movement," as any theory of pragmatics would predict. Participants might still overestimate the probability that Linda is active in the feminist movement, but to consider (B) more likely than (A′) is no longer a violation of elementary probability theory. Accordingly, Tversky and Kahneman (1983) and later work took care to control for this confound and attempt to block the altogether rational but unintended interpretation of option (A). In the original study, controlling for this pragmatic strengthening of option (A) had a mitigating effect in that fewer participants picked the conjunctive option (B) as the most likely, but by no means did mistakes disappear. This prompted Tversky and Kahneman (1983) to conclude that, while interpretive confounds might conspire to create a particularly strong effect in certain conditions, there appears to be an irreducible element of reasoning in this class of problems that does not target the standard notion of probability, and is therefore in need of an explanation.

Similar cases arise in the study of ostensibly deductive problems. Illusory inferences from disjunction, discovered by Walsh and Johnson-Laird (2004), are simple two-premise problems as exemplified below.

(1) John speaks English and Mary speaks French, or Bill speaks German.
John speaks English.
———
Mary speaks French.

Participants draw the conclusion about 85% of the time in behavioral experiments that have been replicated, and independently of whether participants are given the proposed conclusion and asked whether it follows, or instead are given only the two premises and asked to volunteer a conclusion (Walsh & Johnson-Laird, 2004; Koralus & Mascarenhas, 2018; Picat, 2019; Mascarenhas & Koralus, 2016). At first logical blush though, the conclusion does *not* follow: it could be that Bill speaks German (premise one is true), John speaks En-

In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.

glish (premise two is true), but Mary does not speak French (conclusion is false). Classical pragmatic considerations as applied above to the conjunction fallacy won't make a difference here, for the model just described is a counterexample to the inference in (1) even if the "or" in the first premise is interpreted exclusively. Yet, more modern theories of implicature from linguistics do predict an absolving interpretation for premise one. The theory of implicature due to Sauerland (2004), to cite just one example, predicts that premise one of (1) would be interpreted in an *exhaustive* way stronger than simple exclusive disjunction, amounting to "[John speaks English, Mary speaks French, and Bill doesn't speak German] or [Bill speaks German, John doesn't speak English, and Mary doesn't speak French]" (Mascarenhas, 2014).

If this is the interpretation of the first premise of (1), then the conclusion *does* follow from the premises in standard classical logic. Is the illusory inference in (1) no illusion at all then? The facts here are suggestively analogous to those for the conjunction fallacy outlined above. Research in psycholinguistics has shown that the interpretive processes that lead to such exhaustive meanings can be interfered with under cognitive load, increasing the proportion of literal interpretations. Picat (2019) used a dual-task design involving an unrelated memory task to induce cognitive load on reasoners presented with a host of deductive problems, including illusory inferences as in (1). He found that participants were appreciably less likely to draw the proposed fallacious conclusion under cognitive load, precisely for those deductive problems where linguistic theories predict that the "fallacy" is no fallacy at all, and is instead the result of interpretive processes like exhaustification. However, mistakes by no means disappeared, and only a very mild mitigating effect was observed. Much like the case of the conjunction fallacy, it seems that, in illusory inferences from disjunction as in (1), two processes conspire to create a very strong effect: one entirely rational and mistake-free, having to do with implicatures that strengthen the literal meaning in predictable ways, the other more mysterious and in need of an explanation.

There are two related accounts of the non-implicature aspect of these illusory inferences: the mental-model account of Walsh and Johnson-Laird (2004), updated by Khemlani, Byrne, and Johnson-Laird (2018), and the question-answer dynamics of Koralus and Mascarenhas (2013). The two classes of accounts agree on the central driver of the illusion: the operator "or" in the first premise introduces two *alternative possibilities*, and human reasoners use the information in the second premise to draw their attention onto one of those possibilities. Specifically, the two alternative possibilities introduced by the first premise are `John speaks English and Mary speaks French` and `Bill speaks German`, and the second premise `John speaks English`, by virtue of sharing content with the first alternative, causes the second alternative `Bill speaks German` to drop from attention. The reasoner is left with only one alternative `John speaks English and Mary speaks French`, whence the observed conclusion follows straightforwardly.

Koralus and Mascarenhas (2013) cash out this shared central idea in terms of question-answer dynamics. Scholarship from linguistics has pointed out strong connections between disjunction and questions, for example, many entirely unrelated languages of the world share morphemes for the disjunction and question-forming operators (Jayaseelan, 2008; Szabolcsi, Whang, & Zu, 2014). More importantly, linguistic semantics have provided powerful accounts of linguistic phenomena by positing that disjunctions and questions share a fundamental semantic core in that they both introduce alternative possibilities that are available to later linguistic and pragmatic computations (Mascarenhas, 2009; Kratzer & Shimoyama, 2002; Alonso-Ovalle, 2006). Koralus and Mascarenhas (2013) thus propose that the reason humans use information in the second premise of illusory inferences from disjunction to *choose* one of the alternatives provided by the first premise is that they are engaged in a question-answer process. The first premise is akin to a question, and the second premise provides a hint at an answer to that question, as though uttered by an informative and cooperative speaker trying to help the reasoner answer the question at hand. Question-answer dynamics then offers an explanation for the algorithm posited by the original mental-model account, enriching it with entirely independent and well-established observations and formal tools from linguistic semantics.

This entirely informal presentation of the theories is highly incomplete, but it suffices to highlight the questions we investigate in this article. First, we wanted to find out whether non-linguistic materials could trigger the kind of reasoning with alternative possibilities, or question-answer dynamics as we prefer to see it, that is operative in illusory inferences from disjunction. In both the question-answer dynamic account and its earlier mental-model inspiration, the little word "or" and its ability to introduce alternatives are at the core of the phenomenon. But representing questions or alternative possibilities must be an ability of the human mind even in the absence of linguistic stimuli. So we set out to convey the rich structure of (1) purely visually, without linguistic disjunction, and check whether humans still found the fallacious conclusion attractive, suggesting the availability of question-answer dynamics even in the absence of linguistic materials. Second, we tried to shed light on the line between the absolving interpretation of the first premise we discussed earlier in terms of scalar implicature, and the accounts based on processes of reasoning proper. With visually presented premises, scalar implicature is not a plausible mechanism. Thus, if we find any mistakes, they cannot come from an absolving interpretation of the first premise. Conversely, if we find fewer mistakes than in the linguistic version, we can get a better grasp of the extent of the original phenomenon that is due to purely interpretive, linguistic mechanisms. Thirdly, we wanted to contribute toward a paradigm that would allow testing complex deductive problems in non-linguistic populations, that is infants and potentially non-human animals. There is of course much work on these populations' reasoning, but so far it has exclusively focused on simple valid inferences, such as disjunctive syllogism (Cesana-Arlotti et al., 2018; Völter & Call, 2017). In order to advance our understanding of these populations' reasoning faculties, we need to be able to test them with the kind of more complex stimuli that are informative about adult humans' reasoning faculties. The study we report on here makes first steps toward such a paradigm.

Finally, we point out that discovering visual versions of illusory inferences from disjunction, as we do in this study, has ramifications beyond these inferences or even deductive reasoning. Sablé-Meyer and Mascarenhas (2021) have shown that the parallelism between illusory inferences from disjunction and probabilistic reasoning problems from the representativeness literature such as the conjunction fallacy runs much deeper than we've let on in the early paragraphs of this introduction. We cannot do justice to the argument by Sablé-Meyer and Mascarenhas (2021) in this extremely short paper, so we must leave this hopefully tantalizing point unmotivated here, and simply point out that, if question-answer dynamics and alternative possibilities are indeed at the core of the conjunction fallacy as well, then our methodology here has potential applications in the study of visual counterparts of a rather diverse class of attractive fallacies.

Complex deduction from visually presented premises

Design and materials

We investigated illusory inferences from disjunction as in (1), schematized in (2) below.

(2) $(a \wedge b) \vee c$
a
―――――
b

We developed purely visual animations that depict a water-based mechanism. Water flows from the basin at the top through the pipes, and the flow can be blocked by gates that are placed within the pipes. Each closure or opening of the gates represents the truth value of a proposition. As an example, consider the water-based mechanism in Figure 1. There are two passages through which the water can flow from the top to the bottom. If the water were to flow down the pipe, either the left gate or the right gate should be open.

This particular configuration represents the logical structure $a \vee b$ in the following way. After showing that the gates placed within each passage are blocking the water flow, a door covers both gates. Shortly after, the water flows down the pipe and a question mark appears. Given that at least one gate needs to be open for the water to flow down the pipe, this raises the question of whether the left gate opened, or the right gate. Thus, if we assign to the proposition a the truth value of 'the left gate opened' and to the proposition b the truth value of 'the right gate opened', this visual paradigm conveys $a \vee b$ (the left gate opened or the right gate opened).

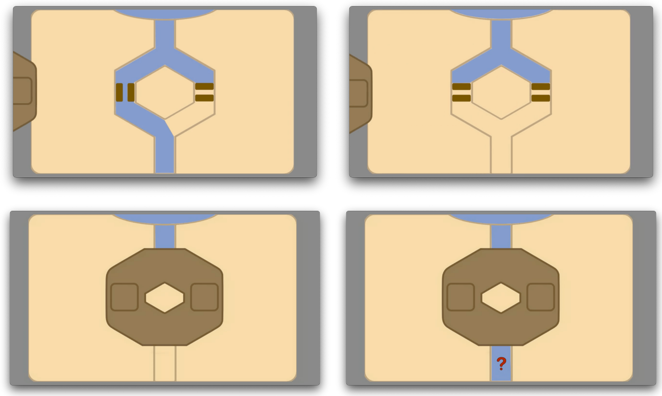

Figure 1: Representing disjunction

We developed two familiarization trials that help the participants get a grasp of how water-based mechanisms work. The first familiarization trial teaches that water flows from the basin at the top to the bottom, and that gates can block the water flow, as depicted in Figure 2.

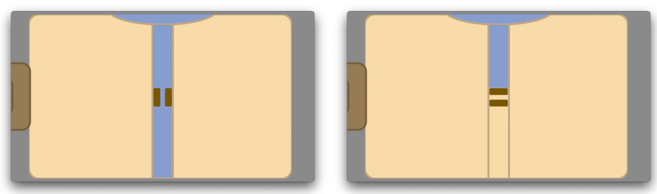

Figure 2: Familiarization trial 1

The second familiarization trial informs that gates can rotate while under cover, and that windows placed within the cover can open to reveal the state of the gates, as illustrated in Figure 3.

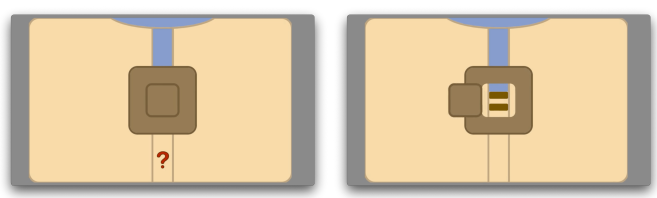

Figure 3: Familiarization trial 2

Following the two familiarization trials, we presented a critical trial, which is either a target trial or a baseline. The target trial—which we call the *Y* scenario—conveys the same content as the inference pattern in (2). As depicted in Figure 4, there are two passages through which the water can flow down the pipe. Two gates were installed within the left pipe and one gate within the right pipe. Initially, the water flow is blocked by the gates in both paths. Then the gates are covered, and shortly after, the water flows down the pipe, raising the question "Did the two gates on the left open or the right gate?" This conveys that $(a \wedge b) \vee c$ is true. We then reveal that the top left gate (i.e., gate A) is open, conveying that *a* is true. Lastly, we ask the participants "Can you conclude that gate B is *guaranteed to be* open?" Logically, one should not be able to conclude that gate B is open because the opening of gate C alone allows the water to flow down the pipe. So if a participant concluded that gate B is guaranteed to be open, then they would be committing a fallacy—precisely the kind attested in the inference pattern in (2).

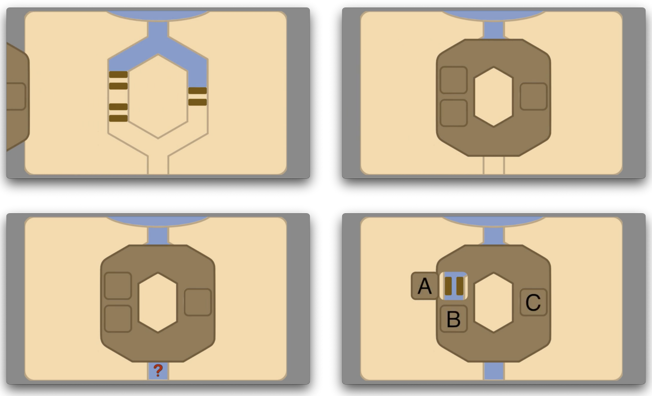

Figure 4: Target trial: the Y scenario

The baseline trial is minimally different from the target trial in that we do not reveal that gate A is open. The corresponding logical structure is provided in (3). It lacks the second premise of the classical illusory inferences from disjunction case represented in (2). Since the second premise is necessary to induce the fallacy, a 'yes' response in the baseline trial is less likely an illusion caused by disjunctive inference. Therefore, by comparing the fallacy rates of the target and baseline trials, we can estimate to what extent the logical structure of illusory inferences from disjunction makes the fallacy attractive. Moreover, by doing so, we expect to filter out biases (if any) arsing from this particular visual setup.

(3) $$\frac{(a \wedge b) \vee c}{b}$$

In addition to the *Y* scenario, we developed another target trial that conveys the same content in order to eliminate potential confounds in the design: The visual proximity of gates A and B in the *Y* scenario could contribute a low-level strategy associating these two. Moreover, the participants' intuitive understanding of fluid dynamics might lead them to think that gate A was opened by the water pressure's being greater on the left side, thereby making it more likely that gate B opened as well.

The visual paradigm depicted in Figure 5—which we call the *Diamond* scenario—aims to address the aforementioned concerns by severing the connection between gates A and B. One gate is installed within each path, and the third gate is placed at the bottom, below the merging point. Initially, all gates are open and the water can flow down the pipe. Then the gates are covered, and the water flows again from the basin at the top. But this time, the water flow is blocked at some point, and raises the question "Did the top two gates close or the bottom gate?" Just as in the Y scenario, the corresponding logical structure is $(a \wedge b) \vee c$. We then reveal that the upper left gate (i.e., gate A) is closed and ask the participants "Can you conclude that gate B is *guaranteed to be* closed?" The logical answer is "no," since the presented scenario is compatible with the situation in which gates A and C are closed but gate B is not. The baseline trial for the Diamond scenario minimally differs from the target trial in that we do not reveal that gate A is closed.

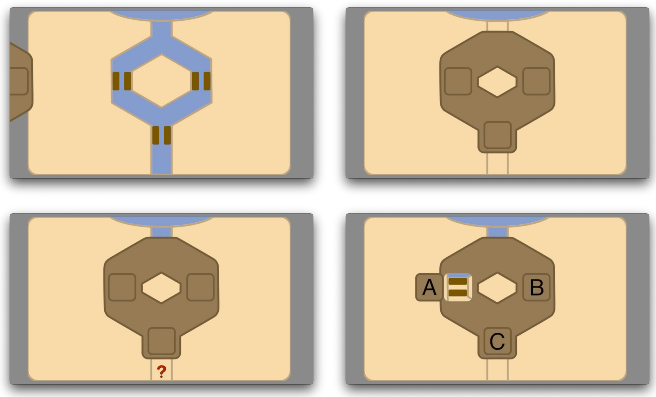

Figure 5: Target trial: the Diamond scenario

The *Diamond* scenario notably differs from the *Y* scenario in two respects. First, in the *Diamond* scenario we ask the participants whether gate B is *closed* whereas in the Y scenario, we ask whether gate B is *open*. Nevertheless, the two scenarios convey the same logical structure due to the ways in which they raise the question $(a \wedge b) \vee c$. In the *Diamond* case, the water flow is blocked and the question concerns which gate is responsible for the blockage. In the *Y* case, all of the gates were initially closed but then the water started to flow down the pipe, and the question concerns which gate is responsible for the flow of water.

Second, gates A and B are not physically connected in the *Diamond* scenario as they are installed within different paths, preventing the aforementioned consideration of fluid dynamics. Moreover, the three gates are equidistant from each other. This likely prevents a low-level strategy associating gates A and B.

Following the target or baseline trial, we presented animations that convey disjunctive syllogism or disjunctive syllogism with an additional disjunct, each of which is illustrated in Figure 6. We used the trials to verify that the participants understood the visual paradigms conveying a simple logical structure. In addition, we used the trial conveying disjunctive syllogism with an additional disjunct (RHS of Figure 6; the *no-control*) as a second baseline, since it should be clear to a cooperative participant that the inference is not valid.

Lastly, the participants were presented with static images

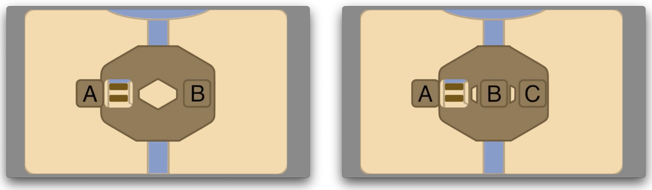

Yes-control
Premise 1: $a \vee b$
Premise 2: $\neg a$
Conclusion: b?

No-control
Premise 1: $a \vee b \vee c$
Premise 2: $\neg a$
Conclusion: b? c?

Figure 6: Control trials

of water-based mechanisms and were asked to assess whether water can flow down the pipe when certain gates were closed. Figure 7 exemplifies such trials. We used the answers to these questions as the means to exclude confused participants.

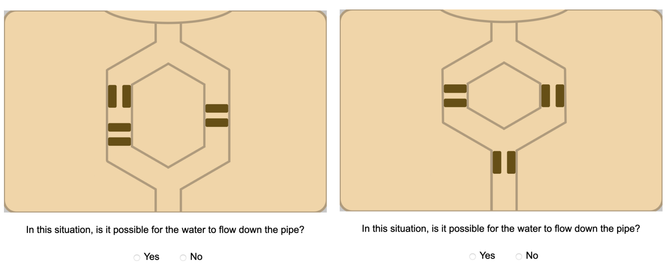

Figure 7: Static controls

Participants and procedure

We recruited 401 participants on Prolific. 64% were female and their mean age was 33. The participants were randomly classified into four groups. All groups were presented with two familiarization trials and three static controls. As for the critical trial, each group was assigned either the *Y* target trial, the *Y* baseline trial, the *Diamond* target trial, or the *Diamond* baseline trial.

Predictions

We expect to observe a higher fallacy rate (i.e., the proportion of yes-responses) in the target trials than in the baseline trials or no-controls. Moreover, since our visual paradigms minimize the effect of language, we predict that the fallacy rate in the target trials to be relatively lower than in the classical illusory inferences from disjunction cases which involve a linguistic disjunction.

We also speculate that the fallacy rate in the *Y* scenario could be higher than in the *Diamond* scenario, because the former is likely to facilitate the development of a low-level strategy associating gates A and B due to their visual proximity and the particpants' intuitive understanding of fluid dynamics. However, we also find the *Diamond* scenario more complex than the *Y* scenario (in the former, neither the gates *A* and *C* nor the gates *B* and *C* form a logical formula despite being physically connected), which could potentially induce more reasoning mistakes across the board.

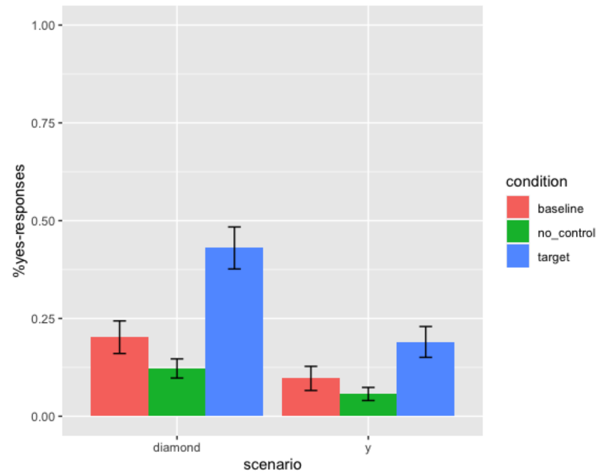

Figure 8: Proportion of yes-responses by scenario and condition; error bars indicate the standard error of the mean. The target conditions of the *Diamond* and *Y* scenarios show 42% and 19% fallacy rates, respectively.

Table 1: Responses to critical trials and no-controls were fitted into a model with SCENARIO, CONDITION, and the interaction between the two as fixed effects, and random intercepts for participants.

Coefficient	**Estimate**	**Standard Error**	***p*-value**
Intercept	0.5509	0.3887	0.1564
Baseline	1.6719	0.5676	0.0032
No-control	2.5563	0.5948	<0.0001
Y	1.9276	0.6155	0.0017
Baseline:Y	-0.8452	0.7939	0.2870
No-control:Y	-0.8998	0.6788	0.1850

Results

We excluded from our analysis the participants who did not properly answer the static control questions. This left us with 373 participants (7% excluded). We analyzed the data using a generalized linear mixed-effects model with the `glmer` function in `R` (Bates, Kliegl, Vasishth, & Baayen, 2015).

We fitted the participants' answers to critical trials and no-controls into a binomial linear mixed-effects model with two predictors: (i) CONDITION with the 3 levels *target/baseline/no-control* (reference level: target) and (ii) SCENARIO with the 2 levels *y/diamond* (reference level: diamond). The largest converging model included random intercepts for participants. Figure 8 plots the proportion of yes-responses by scenario (*Y* vs. *Diamond*) and condition (target vs. baseline vs. no-control) and Table 1 summarizes the fitted model.

We first looked at whether the two scenarios made attrac-

Table 2: Details of the statistical analysis; all comparisons were analyzed using the *glmer* function in *R*, with and random intercepts for participants

Scen.	Comparison	Est.	SE	*p*-value
Dmd.	target vs. baseline	-1.672	0.568	0.0032
Dmd.	target vs. no-ctrl.	-2.556	0.595	<.0001
Dmd.	baseline vs. no-ctr.	-0.884	0.448	0.0481
Y	target vs. baseline	-0.827	0.558	0.1386
Y	target vs. no-ctrl.	-1.657	0.508	0.0011
Y	baseline vs. no-ctrl.	-0.830	0.564	0.1410

tive the fallacious inference akin to (2). We compared the proportion of yes-responses in the target condition to those in the baseline condition and the no-control condition. Our analysis shows that the rate of fallacies ("yes" responses) in the baseline and the no-control were both significantly lower than in the target condition. The rate of fallacies in the *Y* scenario was significantly lower than in the *Diamond* scenario, although there was no significant interaction between the two predictors SCENARIO and CONDITION.

Lastly, we calculated contrasts between conditions using the `emmeans` package in `R`, as summarized in Table 2. We observed that the effect is more pronounced in the *Diamond* scenario than in the *Y* scenario: In the *Diamond* scenario, the differences in target vs. baseline and target vs. no-control were both significant, which is consistent with what we observed in the fitted model. We calculated effect sizes for each contrast, observing medium (Cohen's $W = 0.38$) and high ($W = 0.57$) effect sizes for target vs. baseline and target vs. no-control, respectively. The difference in baseline vs. no-control was marginally significant. Results were mixed for the *Y* scenario. The rate of fallacies in the target condition was significantly higher than in the no-control, and we observed a medium effect size ($W = 0.46$). The difference in target vs. baseline was not significant, although the corresponding effect size was close to medium ($W = 0.24$).

General discussion and conclusions

Our results constitute evidence that it is possible to convey the logical structure in (2) via non-linguistic means. Both in the *Y* scenario and *Diamond* scenario, participants made more mistakes when a purely visual animation conveyed the information in (2) than when an animation represented an inference pattern whose linguistic counterpart is unattractive. These results suggest an important role for interpretive processes like implicature in the linguistic versions of this task, while confirming that the phenomenon cannot be entirely due to linguistic processes. Indeed, it is not difficult to make sense of the lower fallacy rates in our experiments (42% in the *Diamond* scenario and 19% in the *Y* scenario) compared to what was reported for the classical cases of illusory inferences from disjunction (Walsh & Johnson-Laird, 2004, over 80%). Granted that the linguistic interpretive confounds conspire to create a particularly strong effect in certain conditions, we expect a lower endorsement rate when such confounds are eliminated.

Extant reasoning-based theories of the linguistic illusory inference from disjunction identify the source of the inference-making behavior in reasoning processes specific to alternative possibilities. This idea is cashed out in terms of question-answer dynamics by Koralus and Mascarenhas (2013, 2018), incorporating independent insights from linguistics and philosophical logic and providing *a reason* for what in mental-model theory, the classical version or its more recent revision, is an entirely stipulated mechanistic process whereby certain operators like disjunction generate alternative possibilities, and later information is incorporated by dropping some alternatives from consideration. The question-answer framing of the mental-models mechanistic account has proven fruitful uncovering new illusory inferences that share linguistic and logical properties with disjunctions and questions, such as indefinite expressions like "some" (Mascarenhas & Koralus, 2017) and the modal operator "might" (Mascarenhas & Picat, 2019; Bade, Picat, Chung, & Mascarenhas, 2022). In sum, the question-answer approach to these reasoning fallacies isn't merely an extension of the original mental-model theory approach that unifies it with linguistics, it is a properly stronger theory, making novel predictions that do not follow from mental-model theory in any of its incarnations past or present.

If this perspective is on the right track, then question-answer dynamics are pervasive in human reasoning and by no means depend on linguistic sources. But do they depend on *communicative contexts*? While we believe we were successful at factoring out language proper, we did nothing to control for communicative processes: our animations were still of course seen as communicative acts by our participants, intended to impart some content, so that the behavior we observed was plausibly the result of question-answer dynamics being processed online, as the product of properly pragmatic reasoning. Controlling for communicative contexts and reasoning about communicative processes is a daunting task indeed, but the question deserves attention. For the other conceptual possibility is extremely intriguing in our view: question-answer dynamics might be a tool intrinsically available in the human reasoning arsenal, triggered whenever a question is being considered and information is found that could be used to answer it, even absent any trace of communicative intentions.

Acknowledgments

This work was supported by Agence Nationale de la Recherche grants ANR-18-CE28-0008 (LANG-REASON; PI: Mascarenhas) and ANR-17-EURE-0017 (FrontCog; Department of Cognitive Studies, Ecole Normale Supérieure), and by the New Faculty Startup Fund from Seoul National University (Chung).

References

Alonso-Ovalle, L. (2006). *Disjunction in alternative semantics*. Unpublished doctoral dissertation, UMass Amherst.

Bade, N., Picat, L., Chung, W., & Mascarenhas, S. (2022). Alternatives and attention in language and reasoning: a reply to Mascarenhas & Picat (2019). *Semantics & Pragmatics*. doi: 10.3765/sp.15.2

Bates, D., Kliegl, R., Vasishth, S., & Baayen, H. (2015). Parsimonious mixed models. *arXiv preprint arXiv:1506.04967*. doi: 10.48550/arXiv.1506.04967

Cesana-Arlotti, N., Martín, A., Téglás, E., Vorobyova, L., Cetnarski, R., & Bonatti, L. L. (2018). Precursors of logical reasoning in preverbal human infants. *Science*, *359*, 1263–1266. doi: 10.1126/science.aao3539

Jayaseelan, K. (2008). *Question particles and disjunction.* (Unpublished manuscript: The English and Foreign Languages University (Hyderabad))

Khemlani, S., Byrne, R. M. J., & Johnson-Laird, P. N. (2018). Facts and possibilities: a model-based theory of sentential reasoning. *Cognitive Science*, *42*(6), 1887–1924. doi: 10.1111/cogs.12634

Koralus, P., & Mascarenhas, S. (2013). The erotetic theory of reasoning: bridges between formal semantics and the psychology of deductive inference. *Philosophical Perspectives*, *27*, 312–365. doi: 10.1111/phpe.12029

Koralus, P., & Mascarenhas, S. (2018). Illusory inferences in a question-based theory of reasoning. In K. Turner & L. Horn (Eds.), *Pragmatics, truth, and underspecification: Towards an atlas of meaning* (Vol. 34, pp. 300–322). Leiden: Brill. doi: 10.1163/9789004365445_011

Kratzer, A., & Shimoyama, J. (2002). Indeterminate pronouns: the view from Japanese. In *Third tokyo conference on psycholinguistics.*

Mascarenhas, S. (2009). *Inquisitive semantics and logic*. Unpublished master's thesis, ILLC.

Mascarenhas, S. (2014). *Formal semantics and the psychology of reasoning: Building new bridges and investigating interactions*. Unpublished doctoral dissertation, New York University.

Mascarenhas, S., & Koralus, P. (2016). *Free-form response vs. yes/no-question methodologies in the study of human reasoning.* (Poster presented at the 38th Annual Meeting of the Cognitive Science Society)

Mascarenhas, S., & Koralus, P. (2017). Illusory inferences with quantifiers. *Thinking and Reasoning*, *23*(1), 33–48. doi: 10.1080/13546783.2016.1167125

Mascarenhas, S., & Picat, L. (2019). *Might* as a generator of alternatives: the view from reasoning. In *Proceedings of SALT 29, UCLA* (pp. 549–561). doi: 10.3765/salt.v29i0.4635

Picat, L. (2019). *Inferences with disjunction, interpretation or reasoning?* Unpublished master's thesis, Ecole Normale Supérieure.

Sablé-Meyer, M., & Mascarenhas, S. (2021). Indirect illusory inferences from disjunction: a new bridge between deductive inference and representativeness. *Review of Philosophy and Psychology*, *12*(2). doi: 10.1007/s13164-021-00543-8

Sauerland, U. (2004). Scalar implicatures in complex sentences. *Linguistics and Philosophy*, *27*, 367-391. doi: 10.1023/B:LING.0000023378.71748.db

Stenning, K., & van Lambalgen, M. (2008). Interpretation, representation, and deductive reasoning. In J. Adler & L. Rips (Eds.), *Reasoning: studies in human inference and its foundations* (pp. 223–249). Cambridge: Cambridge University Press.

Szabolcsi, A., Whang, J. D., & Zu, V. (2014). Quantifier words and their multi-functional parts. *Language and Linguistics*, *15*(1), 115–155.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: the conjunction fallacy in probability judgment. *Psychological Review*, *90*, 293-315. doi: 10.1037/0033-295X.90.4.293

Völter, C. J., & Call, J. (2017). Causal and inferential reasoning in animals. In J. Call (Ed.), *Apa handbook of comparative psychology: Vol. 2. perception, language, and cognition* (pp. 643–671). The American Psychological Association. doi: 10.1037/0000012-029

Walsh, C., & Johnson-Laird, P. N. (2004). Co-reference and reasoning. *Memory and Cognition*, *32*, 96-106. doi: 10.3758/BF03195823